arXiv:1402.4225v1 [cs.IT] 18 Feb 2014

Information Theory of Matrix Completion

Changho Suh
KAIST, Daejeon, South Korea
Email: chsuh@kaist.ac.kr

***Abstract*—*Matrix completion* is a fundamental problem that comes up in a variety of applications like the Netflix problem, collaborative filtering, computer vision, and crowdsourcing. The goal of the problem is to recover a k-by-n unknown matrix from a subset of its noiseless (or noisy) entries. We define an information-theoretic notion of *completion capacity* C that quantifies the maximum number of entries that one observation of an entry can resolve. This number provides the minimum number m of entries required for reliable reconstruction: $m = \frac{kn}{C}$. Translating the problem into a distributed joint source-channel coding problem with encoder restriction, we characterize the completion capacity for a wide class of *stochastic* models of the unknown matrix and the observation process. Our achievability proof is inspired by that of the Slepian-Wolf theorem. For an arbitrary stochastic matrix, we derive an upper bound on the completion capacity.**

I. Introduction

Matrix completion has received considerable attention with applications in the Netflix problem [1], collaborative filtering [2], computer vision [3], and crowdsourcing [4]. The objective of the problem is to reconstruct a k-by-n unknown matrix from a given subset of noiselessly (or noisily) observed entries. Clearly completely irrelevant entries, if not observed, have no chance to be recovered, since other available entries cannot provide any information about the missing ones. In many applications, however, there can be correlation between entries. For instance, in the Netflix problem, movie ratings of a subscriber, comprising the elements of a single column or row, could play a role to infer ratings of other subscribers having similar tastes in movies. So a natural question that arises in this context is: what is the minimum number of entries needed for reconstruction of the unknown matrix with possibly correlated entries?

There has been a proliferation of reconstruction algorithms [5]–[8] which may suggest an answer. Under the low-rank assumption of the matrix, Candes-Recht [5] made use of the nuclear-norm-based optimization framework to provide a bound on the minimum number of entries required for reconstruction. This bound was further improved by Candes-Tao [6], Keshavan-Montanari-Oh [7], and Recht [8]. However, the best known bound thus far still comes with a numerical constant and/or logarithmic factor gap to the optimal performance, even under the low rank assumption; hence the question raised above has not been answered completely. The challenge might come from the matrix model and the *worst-case* analysis approach that these algorithms have pursued. The algorithms consider a *deterministic* matrix and strive to guarantee reconstruction for *all possible instances* of matrix entries. Moreover, the design of these algorithms is based primarily on computational considerations which might give challenge in finding the fundamental limits that can be achieved without any complexity restriction.

In an effort to make progress towards answering the question, we invoke Shannon's idea in [9]. In the paper, Shannon introduced a *probabilistic* view to communication systems to characterize the maximum amount of information that can be reliably transmitted over a channel. Specifically Shannon viewed the information source and the channel as *random* quantities that can be modeled via probability distributions. Given the stochastic source and channel, he then established the maximum rate of communication as the ratio of the channel capacity to the source entropy rate.

We advocate Shannon's approach for our problem. We model the unknown matrix and the observation process as *stochastic* processes, and seek to find the maximum number of matrix entries that one noiseless (or noisy) observation can reveal on the average. We define the quantity as *completion capacity* C. This number can provide an answer to the question that we raised before. The minimum number m of entries needed for reliable reconstruction of a matrix is the ratio of the total number of matrix entries to the completion capacity: $m = \frac{kn}{C}$.

In this work, we characterize the completion capacity for a wide class of stochastic models of the unknown matrix and the observation process. Drawing parallels between the observation process and the communication channel, we translate the matrix completion problem into a distributed joint source-channel coding problem with encoder restriction. We then invoke the achievability idea of the Slepian-Wolf theorem [10], established for a similar yet different setting, to develop a reconstruction algorithm. We show the optimality of this algorithm for the case in which a sequence of column vectors (or row vectors) of the unknown matrix is a *stationary ergodic* process. This is the only assumption that we make on the statistics of the matrix: we consider *arbitrary* distributions for the entries within each column (or row). The optimality is shown also for arbitrary observation processes which include the *noisy* observation case as well. For an arbitrary stochastic matrix, we develop an upper bound on the completion capacity.

Related Work: The denoising problem [11]–[13] bears similarity to our translated problem in that it intends to estimate an original signal from a noisy version. Especially the problem setting in [13] is intimately related to ours since [13] takes a stochastic view on the original signal and the noisy channel as we do here. However, the problem does not focus on the setting in which only a few entries are observed. Hence the

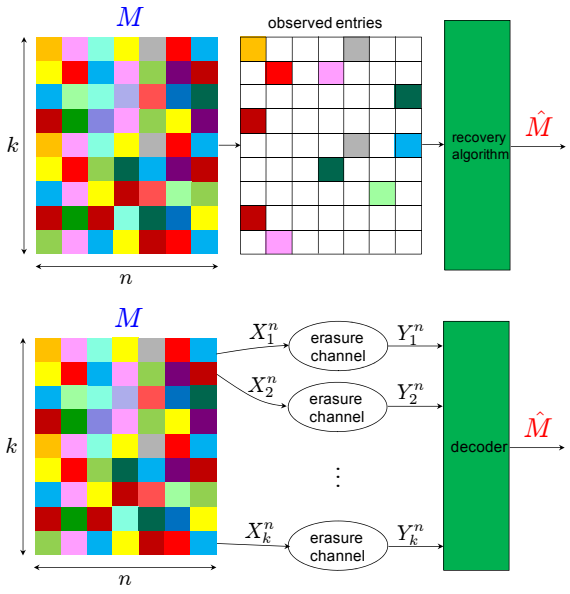

Fig. 1. (Top): The matrix completion problem; (Bottom): A translated joint source-channel coding problem with encoder restriction.

performance metric of our interest, the minimum number of entries for reconstruction, was not investigated therein. On the other hand, Motahari-Bresler-Tse [14] recently applied Shannon's approach that we take here to the DNA sequencing problem. Shannon's idea together with this recent work gave inspiration to our work.

II. Problem Formulation

The problem of matrix completion is illustrated in the top of Fig. 1. Given partially observed entries of a k-by-n unknown matrix M, the task of a recovery algorithm is to decode $\hat{M}$. Here entries are assumed to be sampled uniformly at random. In this work, we translate the problem into a communication problem with encoder restriction. Specifically we model the observation process via an erasure channel with an erasure probability $1-p$, where $p := \frac{m}{kn}$ indicates *observation rate* and m denotes the number of observed entries. Note that the erasure channel is memoryless across entries due to the random sample assumption. Next we view M as an information source that we wish to transmit. The problem can then be regarded as a point-to-point communication problem.

Here one noticeable distinction as compared to the conventional communication setting is that the information source M cannot be processed (encoded) before transmission. In an effort to reflect this encoder restriction, we further reduce the problem into a *distributed* joint source-channel coding problem in which joint processing across different transmitters is not allowed. In this setting, we are forced to somehow respect the encoder constraint, but not fully. Even in this setting, the encoder constraint is still applied to each transmitter.

See the bottom of Fig. 1 for the translated problem. Let X_ℓ^n be the ℓth row (or column) of the unknown matrix M where $\ell \in [1:k] := \{1, \cdots, k\}$. Without loss of generality, assume that X_ℓ^n denotes row components. Here we use shorthand notation to indicate the sequence up to n: $X_\ell^n := (X_{\ell 1}, \cdots, X_{\ell n})$. The (ℓ, i) entry of the matrix, $X_{\ell i}$, is assumed to take values in a finite alphabet set. Since the observation process is modeled via an erasure channel, the channel output w.r.t. the (ℓ, i) entry is

$$Y_{\ell i} = \begin{cases} X_{\ell i}, & \text{w.p. } p; \\ e, & \text{o.w.}, \end{cases} \tag{1}$$

where $\ell \in [1:k]$ and $i \in [1:n]$. Here the observation process is assumed to be noiseless. The noisy observation case will be dealt with in Section V. Decoder uses a function to estimate the unknown matrix M from $\{Y_\ell^n\}_{\ell=1}^k$. An error occurs whenever $\hat{M} \neq M$. The average probability of error is given by $\lambda = \mathbb{E}[\Pr(\hat{M} \neq M)]$.

We say that the *completion rate* $R = \frac{kn}{m} = \frac{1}{p}$ is achievable if there exists a decoder function such that the average decoding error probability of λ goes to zero as n tends to infinity. The *completion capacity* C is the supremum of the achievable completion rates. Note that this notion quantifies the maximum number of matrix entries that one non-erased observation can resolve on the average. With this notion C, the minimum number of entries required to reconstruct a k-by-n matrix can be expressed as $\frac{kn}{C}$.

Remark 1: Since rows and columns are interchangeable with a slight modification of notation, the completion capacity can be defined in the limit of k alternatively. This asymptotic-regime analysis can be useful in practice as well since in many applications the matrix dimension admits a high dimensional setting, e.g., in the Netflix problem [1], $(k, n) \approx (5 \cdot 10^5, 2 \cdot 10^4)$. □

Remark 2: Our problem setting with multiple transmitters comes with another benefit. Since it matches the two-dimensional structure of the matrix, it helps to well capture the stochastic property of the matrix. In general, the statistics of the unknown matrix is governed by an arbitrary joint distribution. For some applications, however, the matrix may possess a column (or row) dependent stochastic property. For instance, consider a setting in the Netflix problem where the column index indicates movies of *different* genres and the row index denotes subscribers with *similar* tastes. The different genres may render movie ratings less correlated across columns. In the extreme case where the ratings are independent, we may have $p(x_{11}, \cdots, x_{kn}) \approx \prod_{i=1}^n p(x_{1i}, \cdots, x_{ki})$. On the other hand, ratings across rows may be strongly dependent as they are affected by similar tastes of subscribers. In an extreme case, $p(x_{\ell 1}, \cdots, x_{\ell n})$ may be the same for all $\ell \in [1:k]$. The two-dimensional structure of our setting enables us to reflect this stochastic structure in a convenient manner. It turns out this makes analysis more tractable and simpler rather than in the point-to-point setting. This will be clearer in the subsequent sections. □

III. I.I.D. Process Model

We first consider a simple stochastic model in which a sequence of column vectors $\{(X_{1i}, \cdots, X_{ki})\}_{i=1}^n$ is indepen-

dent and identically distributed (i.i.d.) over i. Notice that this model represents the extreme case in Remark 2, to some extent, where the movie scores are completely irrelevant across different genres. Here the entries within each column are assumed to follow an arbitrary distribution. Obviously this model is very simplistic and hence it may be far from the statistics of a realistic matrix in practical applications. But the i.i.d. model turns out to lay the foundation to tackle more complicated yet practically-relevant cases. Moreover it can serve as a lower bound to the capacity for the general case in which a sequence of column vectors are possibly dependent. This will be clearer in Sections IV and V.

Theorem 1 (I.I.D. Model):

$$C = \frac{\sum_{\ell=1}^{k} H(X_\ell)}{H(X_1, \cdots, X_k)}. \tag{2}$$

Proof: See the following two subsections. ■

Remark 3: For the case where the entries of a column vector are also i.i.d., i.e., $p(x_1, \cdots, x_k) = \prod_{\ell=1} p(x_\ell)$, we get: $C = 1$. This coincides with our intuition because in that case one observation cannot provide any information about others. For the other extreme case where $X_1 = \cdots = X_k$, we have $C = k$. This result also makes sense since in that case one observation can reveal other $k-1$ entries in the same column. □

A. Proof of Achievability

Our proof bears a resemblance to that of the Slepian-Wolf theorem [10], developed in the context of a similar setting. The achievable scheme in [10] employs random binning and *joint typicality decoding*. Encoder ℓ independently partitions the space (that the sequence X_ℓ^n resides in) into multiple bins depending on a rate assigned. It then sends the bin index to which X_ℓ^n belongs. Upon receiving the bin indices from k encoders, the decoder looks for a sequence tuple $(X_1^n, \cdots, X_k^n)$ such that their corresponding bin indices coincide with the received ones and the sequence tuple is jointly typical.

However, this scheme is not directly applicable to our case, since in our model 1) there is no explicit encoder; and 2) X_ℓ^n's are observed through erasure channels instead of through noiseless bit-pipes. These differences prevent employing the binning scheme. So the decoding rule needs to be modified accordingly. The key observation for the modification is that by the asymptotic equipartition property (AEP) of an i.i.d. process, X_ℓ^n is a typical sequence for sufficiently large n, and so is (X_ℓ^n, Y_ℓ^n) [15], [16]. This leads us to consider the following decoding rule. For illustration purpose, we first focus on the case of $k = 2$. The general case of an arbitrary value of k will be dealt with later. Given (y_1^n, y_2^n), the decoder tries to find a pair of sequences $(\hat{x}_1^n, \hat{x}_2^n)$ such that $(\hat{x}_1^n, y_1^n) \in A_\epsilon^{(n)}(X_1, Y_1)$, $(\hat{x}_2^n, y_2^n) \in A_\epsilon^{(n)}(X_2, Y_2)$ and $(\hat{x}_1^n, \hat{x}_2^n) \in A_\epsilon^{(n)}(X_1, X_2)$, where $A_\epsilon^{(n)}(X_1, Y_1)$ denotes a joint typical set w.r.t. X_1^n and Y_1^n for some $\epsilon > 0$. Similarly $A_\epsilon^{(n)}(X_2, Y_2)$ and $A_\epsilon^{(n)}(X_1, X_2)$ are defined accordingly. For notational simplicity, we will use a common notation $A_\epsilon^{(n)}$, which can be easily differentiated from contexts.

Now let us consider the probability of error. Using the union bound, we can upper-bound the probability of error as:

$$\begin{aligned}\Pr\left\{(\hat{X}_1^n, \hat{X}_2^n) \neq (X_1^n, X_2^n)\right\} &\leq \Pr\left\{(X_1^n, X_2^n) \notin A_\epsilon^{(n)}\right\} \\ &+ \sum_{(x_1^n, x_2^n) \in A_\epsilon^{(n)}} p(x_1^n, x_2^n) \left\{\Pr(E_1) + \Pr(E_2) + \Pr(E_{12})\right\}\end{aligned}$$

where E_ℓ indicates the event that $\hat{X}_\ell^n \neq x_\ell^n$, and E_{12} denotes the event that $\hat{X}_1^n \neq x_1^n$ and $\hat{X}_2^n \neq x_2^n$. By AEP, $\Pr\{(X_1^n, X_2^n) \notin A_\epsilon^{(n)}\} \to 0$ as n tends to infinity. Now consider

$$\begin{aligned}\Pr(E_1) &= \sum_{\hat{x}_1^n \neq x_1^n, \hat{x}_1^n \in A_\epsilon^{(n)}(X_1|x_2^n)} \Pr\left\{(\hat{x}_1^n, y_1^n) \in A_\epsilon^{(n)}\right\} \\ &\leq 2^{-n(I(X_1;Y_1) - H(X_1|X_2) - 2\epsilon)}\end{aligned}$$

where the inequality follows from $|A_\epsilon^{(n)}(X_1|x_2^n)| \leq 2^{n(H(X_1|X_2)+\epsilon)}$ for $x_2^n \in A_\epsilon^{(n)}$ [15], [16]. Similarly we get: $\Pr(E_2) \leq 2^{-n(I(X_2;Y_2) - H(X_2|X_1) - 2\epsilon)}$ and $\Pr(E_{12}) \leq 2^{-n(I(X_1;Y_1) + I(X_2;Y_2) - H(X_1,X_2) - 3\epsilon)}$. Hence, the probability of error can be made arbitrarily close to zero (as n tends to infinity) if $I(X_1;Y_1) \geq H(X_1|X_2)$, $I(X_2;Y_2) \geq H(X_2|X_1)$, and $I(X_1;Y_1) + I(X_2;Y_2) \geq H(X_1, X_2)$. Note that $I(X_\ell;Y_\ell) = H(X_\ell) - H(X_\ell|Y_\ell) = pH(X_\ell)$. Applying this to the above condition, we get: $p \geq \max\left\{\frac{H(X_1|X_2)}{H(X_1)}, \frac{H(X_2|X_1)}{H(X_2)}, \frac{H(X_1,X_2)}{H(X_1)+H(X_2)}\right\}$. Using the non-negativity of entropy and the fact that conditioning reduces entropy, we get $p \geq \frac{H(X_1,X_2)}{H(X_1)+H(X_2)}$. Since $R = \frac{1}{p}$, this condition becomes $R \leq \frac{H(X_1)+H(X_2)}{H(X_1,X_2)}$, which proves the achievablility for $k = 2$.

Precisely similar arguments can be made for an arbitrary value of k. Specifically, given $(y_1^n, \cdots, y_k^n)$, the decoder looks for a tuple of k sequences $(\hat{x}_1^n, \cdots, \hat{x}_k^n)$ such that $(\hat{x}_\ell^n, y_\ell^n) \in A_\epsilon^{(n)}, \forall \ell \in [1:k]$ and $(\hat{x}_1^n, \cdots, \hat{x}_k^n) \in A_\epsilon^{(n)}$. Then, the probability of error can be made arbitrarily close to zero if $\sum_{\ell \in \mathcal{S}} I(X_\ell;Y_\ell) \geq H(X_\mathcal{S}|X_{\mathcal{S}^c}), \forall \mathcal{S} \subseteq [1:k]$, where $X_\mathcal{S} = \{X_\ell\}_{\ell \in \mathcal{S}}$. This yields the desired result: $R \leq \min_{\mathcal{S} \subseteq [1:k]} \frac{\sum_{\ell \in \mathcal{S}} H(X_\ell)}{H(X_\mathcal{S}|X_{\mathcal{S}^c})} = \frac{\sum_{\ell=1}^k H(X_\ell)}{H(X_1, \cdots, X_k)}$.

B. Proof of Converse

Using $I(X_\ell;Y_\ell) = pH(X_\ell)$ and Fano's inequality $H(X_1^n, \cdots, X_k^n | Y_1^n, \cdots, Y_k^n) \leq n\epsilon_n$, we get:

$$\begin{aligned}\frac{n\sum_{\ell=1}^k H(X_\ell)}{R} &= \sum nI(X_\ell;Y_\ell) \overset{(a)}{=} \sum I(X_\ell^n;Y_\ell^n) \\ &\overset{(b)}{\geq} H\left(Y_1^n, \cdots, Y_k^n\right) - \sum H(Y_\ell^n|X_\ell^n) \\ &\overset{(c)}{=} I\left(X_1^n, \cdots, X_k^n; Y_1^n, \cdots, Y_k^n\right) \\ &\overset{(d)}{\geq} nH(X_1, \cdots, X_k) - n\epsilon_n\end{aligned}$$

where (a) follows from the memoryless property of the channel and the i.i.d. assumption on X_ℓ^n; (b) follows from the fact that conditioning reduces entropy; (c) follows from a Markov chain of $\{X_j^n, Y_j^n\}_{j \neq \ell} \to X_\ell^n \to Y_\ell^n$; and (d) follows

from Fano's inequality. If R is achievable, then $\epsilon_n \to 0$ as n tends to infinity. Hence, we obtain the desired bound, as claimed.

IV. Stationary Ergodic Process Model

We extend the result for the i.i.d. case to a more general case in which a sequence of column vectors $\{(X_{1i}, \cdots, X_{ki})\}_{i=1}^{\infty}$ is now a stationary ergodic process. As before, the entries within each column are assumed to follow an arbitrary distribution.

Theorem 2 (Stationary Ergodic Model):

$$C = \frac{\sum_{\ell=1}^{k} \bar{H}(X_\ell) - \sum_{\ell=1}^{k} a_\ell}{\bar{H}(X_1, \cdots, X_k) - \sum_{\ell=1}^{k} a_\ell} \tag{3}$$

where $\bar{H}(X_\ell) := \lim_{n\to\infty} \frac{1}{n} H(X_\ell^n)$ indicates the entropy rate, $\bar{H}(X_1, \cdots, X_k) := \lim_{n\to\infty} \frac{1}{n} H(X_1^n, \cdots, X_k^n)$, and $a_\ell := \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} I(Y_{\ell,(i+1)}^n; X_{\ell i} | X_\ell^{i-1})$.

Proof: See the following two subsections. ■

Remark 4: Unlike the i.i.d. case, here we have a_ℓ. One can interpret this quantity as the information that the present provides about the future noisy signals knowing the past. So this can be alternatively expressed via directed information [17]. In the stationary ergodic case, correlation between entries can make a_ℓ strictly positive, and this contributes to increasing completion capacity. This result makes sense since with more correlation, obviously more entries can be revealed per observation. Also note that this theorem includes the result for the i.i.d. model as a special case. Specializing to the i.i.d. model, we get: $a_\ell = 0$, $\bar{H}(X_\ell) = H(X_\ell)$, $\forall \ell$ and $\bar{H}(X_1, \cdots, X_k) = H(X_1, \cdots, X_k)$, thus yielding (2). In general, $a_\ell \geq 0$, $\bar{H}(X_\ell) \leq H(X_\ell)$, and $\bar{H}(X_1, \cdots, X_k) \leq H(X_1, \cdots, X_k)$. From this, we see that the capacity in the i.i.d. case serves as a lower bound, as claimed earlier. □

A. Proof of Achievability

The proof relies primarily on the Shannon-McMillan-Breiman theorem [9], [18], [19], which showed that for a stationary ergodic process, the AEP holds as in the i.i.d. case: $-\frac{1}{n} \log p(X_1^n, \cdots, X_k^n) \to \bar{H}(X_1, \cdots, X_k)$ and $-\frac{1}{n} \log p(X_\ell^n) \to \bar{H}(X_\ell)$ with probability 1. Using the stationarity of $\{X_{\ell i}\}_{i=1}^{\infty}$ and the memoryless property of the channel, one can show that $\{X_{\ell i}, Y_{\ell i}\}_{i=1}^{\infty}$ is also stationary and ergodic $\forall \ell \in [1:k]$. Moreover, $\{Y_{\ell i}\}_{i=1}^{\infty}$ is necessarily stationary and ergodic since it is a projection of $\{X_{\ell i}, Y_{\ell i}\}_{i=1}^{\infty}$. Hence, the AEP holds for these random processes as well. This naturally motivates us to apply the same decoding rule that we used in the i.i.d. case. Given $(y_1^n, \cdots, y_k^n)$, the decoder looks for a unique tuple of k sequences $(\hat{x}_1^n, \cdots, \hat{x}_k^n)$ such that $(\hat{x}_\ell^n, y_\ell^n) \in \bar{A}_\epsilon^{(n)}, \ell \in [1:k]$ and $(\hat{x}_1^n, \cdots, \hat{x}_k^n) \in \bar{A}_\epsilon^{(n)}$. Here $\bar{A}_\epsilon^{(n)}$ indicates a joint typical set w.r.t. some stationary ergodic processes. One can then readily show that the probability of error goes to zero if $\sum_{\ell\in\mathcal{S}} \bar{I}(X_\ell; Y_\ell) \geq \bar{H}(X_\mathcal{S} | X_{\mathcal{S}^c}), \forall \mathcal{S} \subseteq [1:k]$, where $\bar{I}(X_\ell; Y_\ell) := \lim_{n\to\infty} \frac{1}{n} I(X_\ell^n; Y_\ell^n)$ denotes the mutual information rate.

Unlike the i.i.d. case, in the stationary ergodic case, $\bar{I}(X_\ell; Y_\ell) \neq p\bar{H}(X_\ell)$ in general. Actually we find the relation between $p = \frac{1}{R}$ and $\bar{I}(X_\ell; Y_\ell)$ from the following computation:

$$\begin{aligned} I(X_\ell^n; Y_\ell^n) &= H(X_\ell^n) - \sum H(X_{\ell i} | X_\ell^{i-1}, Y_\ell^n) \\ &= H(X_\ell^n) - \sum H(X_{\ell i} | X_\ell^{i-1}, Y_{\ell i}, Y_{\ell,(i+1)}^n) \\ &= H(X_\ell^n) - (1-p) \sum H(X_{\ell i} | X_\ell^{i-1}, Y_{\ell,(i+1)}^n) \\ &= pH(X_\ell^n) + (1-p) \sum I(X_{\ell i}; Y_{\ell,(i+1)}^n | X_\ell^{i-1}) \end{aligned} \tag{4}$$

where the second equality is due to a Markov chain of $Y_\ell^{i-1} \to X_\ell^{i-1} \to X_{\ell i}$. Notice that $a_\ell := \lim_{n\to\infty} \frac{1}{n} \sum I(X_{\ell i}; Y_{\ell,(i+1)}^n | X_\ell^{i-1})$ *exists* due to the stationarity of (X_ℓ^n, Y_ℓ^n) and X_ℓ^n. Using the above, we then get: $R \leq \min_{\mathcal{S}\subseteq[1:k]} \frac{\sum_{\ell\in\mathcal{S}} \{\bar{H}(X_\ell) - a_\ell\}}{\bar{H}(X_\mathcal{S} | X_{\mathcal{S}^c}) - \sum_{\ell\in\mathcal{S}} a_\ell}$. Using $\bar{H}(X_\ell) - a_\ell \geq 0$ and the fact that conditioning reduces entropy, we finally obtain the desired result.

B. Proof of Converse

Using (4) and making similar arguments as in the i.i.d. case, we get:

$$\begin{aligned} \frac{\sum_{\ell=1}^{k} \{\bar{H}(X_\ell) - a_\ell\}}{R} + \sum a_\ell &= \sum \bar{I}(X_\ell; Y_\ell) \\ &\geq \lim_{n\to\infty} \frac{1}{n} \left[H(Y_1^n, \cdots, Y_k^n) - \sum H(Y_\ell^n | X_\ell^n) \right] \\ &= \lim_{n\to\infty} \frac{1}{n} I(X_1^n, \cdots, X_k^n; Y_1^n, \cdots, Y_k^n) \\ &\geq \bar{H}(X_1, \cdots, X_k) - \lim_{n\to\infty} \epsilon_n. \end{aligned}$$

This yields the desired bound.

V. Generalization

We extend the noiseless observation case to the noisy case illustrated in Fig. 2. We model the noisy observation via a discrete memoryless channel described by a conditional probability distribution. The noisy version Y_ℓ^n then passes through an erasure channel to produce Z_ℓ^n, $\forall \ell$. We will first characterize the completion capacities for the i.i.d. model and for the stationary ergodic model. We will then derive an upper bound for an arbitrary stochastic matrix.

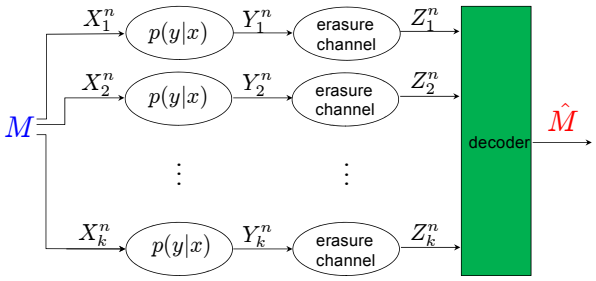

Fig. 2. Noisy observation process.

Theorem 3 (Noisy Entries): For the i.i.d. process model,

$$C = \frac{\sum_{\ell=1}^{k} I(X_\ell; Y_\ell)}{H(X_1, \cdots, X_k)}. \tag{5}$$

For the stationary ergodic process model,

$$C = \frac{\sum_{\ell=1}^{k} \bar{I}(X_\ell; Y_\ell) - \sum_{\ell=1}^{k} b_\ell}{\bar{H}(X_1, \cdots, X_k) - \sum_{\ell=1}^{k} b_\ell} \tag{6}$$

where $b_\ell := \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} I(Y_{\ell i}; Z_{\ell,(i+1)}^n | Y_\ell^{i-1})$.

Remark 5: Observe that the entropy and the entropy rate that appeared in the noiseless case ((2) and (3)) are now replaced by the mutual information and the mutual information rate respectively. This clearly shows the reduction in capacity due to noises. □

Proof: The proof for the i.i.d. case is straightforward. The key observation here is that (X_ℓ^n, Z_ℓ^n) is also i.i.d. This leads us to apply the same decoding rule as before. One can then easily obtain the following condition for achievability: $\sum_{\ell\in\mathcal{S}} I(X_\ell; Z_\ell) \geq H(X_\mathcal{S}|X_{\mathcal{S}^c}), \forall \mathcal{S} \subseteq [1:k]$. We now consider: $I(X_\ell; Z_\ell) = H(X_\ell) - (1-p)H(X_\ell) - pH(X_\ell|Y_\ell) = pI(X_\ell; Y_\ell)$. Applying this to the above condition, we can readily prove the achievability. In the i.i.d. case, $I(X_\ell^n; Z_\ell^n) = pI(X_\ell^n; Y_\ell^n)$. So we get: $\frac{\sum_{\ell=1}^{k} I(X_\ell^n; Y_\ell^n)}{R} = \sum I(X_\ell^n; Z_\ell^n) \geq nH(X_1, \cdots, X_k) - n\epsilon_n$. This establishes the converse proof.

Unlike the i.i.d. case, for the stationary ergodic model, $I(X_\ell^n; Z_\ell^n) \neq pI(X_\ell^n; Y_\ell^n)$ in general. Instead we get:

$$\begin{aligned}
&I(X_\ell^n; Z_\ell^n) = I(X_\ell^n; Y_\ell^n) - I(X_\ell^n; Y_\ell^n | Z_\ell^n) \\
&= I(X_\ell^n; Y_\ell^n) - \sum (1-p) I(X_\ell^n; Y_{\ell i} | Y_\ell^{i-1}, Z_{\ell,(i+1)}^n) \\
&= pI(X_\ell^n; Y_\ell^n) + \sum (1-p) \cdot \\
&\qquad \{I(X_\ell^n; Y_{\ell i} | Y_\ell^{i-1}) - I(X_\ell^n; Y_{\ell i} | Y_\ell^{i-1}, Z_{\ell,(i+1)}^n)\} \\
&= pI(X_\ell^n; Y_\ell^n) + \sum (1-p) I(Y_{\ell i}; Z_{\ell,(i+1)}^n | Y_\ell^{i-1})
\end{aligned} \tag{7}$$

where the first step follows from $X_\ell^n \to Y_\ell^n \to Z_\ell^n$; and the last step is due to $(Y_\ell^{i-1}, Z_{\ell,(i+1)}^n) \to X_\ell^n \to Y_{\ell i}$. From (7), we get: $\bar{I}(X_\ell; Z_\ell) = p(\bar{I}(X_\ell; Y_\ell) - b_\ell) + b_\ell$. Using this and making the same arguments as before, one can readily prove the achievability and the converse. ■

Theorem 4 (Arbitrary Stochastic Matrix):

$$C \leq \limsup_{n\to\infty} \frac{\sum_\ell I(X_\ell^n; Y_\ell^n) - \sum_\ell b_\ell(n)}{H(X_1^n, \cdots, X_k^n) - \sum_\ell (b_\ell(n) - c_\ell(n))}, \tag{8}$$

where $b_\ell(n) := \sum_{i=1}^{n} I(Y_{\ell i}; Z_{\ell,(i+1)}^n | Y_\ell^{i-1})$ and $c_\ell(n) := I(Z_\ell^n; Z_1^n, \cdots, Z_{\ell-1}^n)$.

Proof: Starting with (7), we get:

$$\begin{aligned}
&\frac{\sum \{I(X_\ell^n; Y_\ell^n) - b_\ell(n)\}}{R} + \sum b_\ell(n) = \sum I(X_\ell^n; Z_\ell^n) \\
&= \left[H(Z_1^n, \cdots, Z_k^n) - \sum H(Z_\ell^n | X_\ell^n) \right] + \sum c_\ell(n) \\
&= I(X_1^n, \cdots, X_k^n; Z_1^n, \cdots, Z_k^n) + \sum c_\ell(n) \\
&\geq H(X_1^n, \cdots, X_k^n) + \sum c_\ell(n) - n\epsilon_n.
\end{aligned}$$

This completes the proof. ■

VI. Conclusion

We established the completion capacity for a class of stochastic matrices, and developed an upper bound for arbitrary matrices. Our results can be used in evaluating different reconstruction algorithms, as well as provide the limits that an optimal algorithm can achieve. This paper comes with some limitations. The stationary ergodic model for the unknown matrix mainly considered herein may be far from the statistics of a realistic matrix in some important applications. Moreover, our reconstruction algorithm is based on joint typicality decoding which faces challenge in implementation. Hence future works of interest would be (1) characterizing the completion capacity for an arbitrary stochastic matrix; (2) developing an optimal low-complexity algorithm that achieves the limits possibly without knowing the statistics of the matrix.

Acknowledgment

The author would like to thank Prof. Jinwoo Shin for discussions in the early stage of this work.

References

[1] J. Bennett and S. Lanning, "The Netflix Prize," *Proceedings of KDD Cup and Workshop*, Aug. 2007.

[2] D. Goldberg, D. Nichols, B. M. Oki, and D. Terry, "Using collaborative filtering to weave an information tapestry," *Communications of the ACM*, vol. 35, pp. 61–70, Dec. 1992.

[3] P. Chen and D. Suter, "Recovering the missing components in a large noisy low-rank matrix: Application to SFM source," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, pp. 1051–1063, Aug. 2004.

[4] D. R. Karger, S. Oh, and D. Shah, "Budget-optimal task allocation for reliable crowdsourcing systems," *arix:1110.3564*, Oct. 2011.

[5] E. Candes and B. Recht, "Exact matrix completion via convex optimization," *Foundations of Computational Mathematics*, vol. 9, pp. 717–772, June 2009.

[6] E. J. Candes and T. Tao, "The power of convex relaxation: Near-optimal matrix completion," *IEEE Transactions on Information Theory*, vol. 56, pp. 2053–2080, May 2010.

[7] R. H. Keshavan, A. Montanari, and S. Oh, "Matrix completion from a few entries," *IEEE Transactions on Information Theory*, vol. 56, pp. 2980–2998, June 2010.

[8] B. Recht, "A simpler approach to matrix completion," *Journal of Machine Learning Research*, vol. 12, pp. 3413–3430, 2011.

[9] C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, pp. 379–423, July 1948.

[10] D. Slepian and J. Wolf, "Noiseless coding of correlated information sources," *IEEE Transactions on Information Theory*, vol. 19, pp. 471–480, July 1973.

[11] D. L. Donoho, "De-noising by soft-thresholding," *IEEE Transactions on Information Theory*, vol. 41, pp. 613–627, May 1995.

[12] D. L. Donoho, I. M. Johnstone, G. Keryacharian, and D. Picard, "Wavelet shrinkage: Asymptopia?," *Journal of the Royal Statstical Society*, vol. 57, pp. 301–369, 1995.

[13] T. Weissman, E. Ordentlich, G. Seroussi, S. Verdu, and M. J. Weinberger, "Univeral discrete denoising: Known channel," *IEEE Transactions on Information Theory*, vol. 51, pp. 5–28, Jan. 2005.

[14] A. S. Motahari, G. Bresler, and D. N. C. Tse, "Information theory of DNA shotgun sequencing," *IEEE Transactions on Information Theory*, vol. 59, pp. 6273–6289, Oct. 2013.

[15] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York Wiley, 2th ed., July 2006.

[16] A. E. Gamal and Y.-H. Kim, *Network Information Theory*. Cambridge University Press, 2011.

[17] J. L. Massey, "Causality, feedback, and directed information," *Proceedings of International Symposium on Information Theory and its Applications*, pp. 303–305, 1990.

[18] B. McMillan, "The basic theorems of information theory," *The Annals of Mathematical Statistics*, vol. 24, pp. 196–219, Feb. 1953.

[19] L. Breiman, "The individual ergodic theorem of information theory," *The Annals of Mathematical Statistics*, vol. 28, pp. 809–810, 1960.